Space Operations Communicator
Vol. 5, No. 4, October – December 2008

An Interplanetary Internet

Vinton G. Cerf[1,2]
Jet Propulsion Laboratory, California Institute of Technology
Pasadena, CA 91109

Manned and robotic space exploration in the future will require more advanced and richer networking options than have heretofore been available. This brief essay offers perspective on desirable features of such a network architecture and a high level view of the current status of a decade-long project aimed at the design, implementation, and space-qualification of an interplanetary networking architecture that could satisfy space communication requirements during the next several decades.

Nomenclature

AOS	=	Advanced Orbiting Systems
BP	=	Bundle Protocol [see RFC 5050]
CFDP	=	CCSDS File Delivery Protocol
CCSDS	=	Consultative Committee on Space Data Systems
DARPA	=	Defense Advanced Research Projects Agency
DNS	=	Domain Name System (of the Internet)
DTN	=	Delay and Disruption Tolerant Networking
DTNRG	=	DTN Research Group [part of the IRTF]
IETF	=	Internet Engineering Task Force
IP	=	Internet Protocol
IRTF	=	Internet Research Task Force
LTP	=	Licklider Transport Protocol
RFC	=	Request for Comment [see www.rfc-editor.org]
TCP/IP	=	Transmission Control Protocol/Internet Protocol
URI	=	Uniform Resource Identifier
URL	=	Uniform Resource Locator
URN	=	Uniform Resource Name

I. Introduction

THE exploration of the solar system using robotic technology and manned exploration of near-Earth space has been an ongoing activity since the 1960s. The Mariner mission series launched the first US interplanetary explorations to Mars and Venus, for example. In the ensuing 45 years, the Jet Propulsion Laboratory's Deep Space Network (re-named the Interplanetary Network in recent years) has provided basic communication for a substantial portion of all deep space exploration during this period. For the most part, interplanetary missions have used point-to-point radio links to communicate between Earth and the spacecraft. In some instances, radio relays have been used to extend range or to serve where line-of-sight methods were not feasible. In 1998, a small team of engineers at JPL and MITRE began to consider the possibility of an interplanetary extension of the Internet that they called "InterPlaNet" or IPN (see Acknowledgment section).

The Interplanetary Internet was envisioned as a rich networking infrastructure that could provide for more elaborate mission communication architectures including orbiting and surface resources in manned and robotic

---

[1] Distinguished Visiting Scientist, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109

[2] Cerf is also Vice President and Chief Internet Evangelist at Google, inc. The research described in this paper, carried out at the Jet Propulsion Laboratory, California Institute of Technology, was done so under a contract with the National Aeronautics and Space Administration.

configurations. As an exercise in long-range thinking, the team tried to imagine possibilities on the order of a hundred years in the future when permanent orbiting and landed resources might be in regular communication and when local communication might be as important as linking space-based assets with Earth. Early in the process, it was concluded that simply deploying the current Internet protocols would not be adequate to cope with the high delays caused by speed of light propagation over interplanetary distances and the disruptions produced by planetary rotation, satellite orbital dynamics, and the effects of solar wind or radiation storms on radio communication.

The Internet suite of protocols (TCP, IP, DNS, IP routing protocols, etc.) were conceived and designed to operate in relatively low delay (hundreds of milliseconds) environments and in reliably connected conditions. These protocols have embedded expectations of response times that would not readily adapt to the forty minute round trip times encountered for Earth-Mars communication when the planets are farthest apart. Of course, one might tune particular implementations of the Internet protocols for this purpose but then all communicating systems on the terrestrial Internet would then need to be similarly "tuned" if they were to communicate successfully with off-planet resources. The problem is exacerbated by the possibility that terrestrial devices on the Internet may be mobile and may change IP addresses as a consequence of mobility. An off-planet correspondent using the conventional Domain Name System to look up IP addresses might well use an out-of-date address and fail to reach the mobile destination. While there are potential work-arounds to deal with these and other problems, the team chose to develop a suite of protocols specifically designed to work in high, variable delay conditions in which connectivity might be frequently disrupted. Exploring these ideas led to the invention of Delay and Disruption Tolerant Networking (DTN) protocols[1].

II. Delay and Disruption Tolerant Networking

As the team pursued alternative communication architectures, it became clear that the interplanetary network ideas were a special case of a more general problem of reliable communication in highly disrupted and variably delayed environments. As an example, nodes of sensor networks typically have their communications subsystems powered off most of the time to conserve power. This manifests itself as a disruption in communications with the node. If a message has to traverse a multi-hop field of such sensors, it might take considerable time before the right combinations of sensors power on their communications systems and forward the message, resulting in a long delivery latency. These considerations led the team to propose the creation of an Internet Research Task Force (IRTF) research group on delay and disruption tolerant networking (DTNRG) and the publication of two specifications in the form of Requests for Comment (RFCs). The first of these described the basic architecture of such networks[2] and a second detailed the specification of the Bundle Protocol[3]. Other specifications are in Internet-Draft form for discussion by the working group.

The Bundle Protocol is organized around the idea that "bundles" of data would be packaged up and sent from a source to a destination, possibly in some other part of the solar system. However, it was considered reasonable to consider local interactions by means of bundles as well. In some scenarios, all devices exchanging bundles might be nearby but possibly not in line of sight or perhaps operating in a low power mode that precludes receipt of data or response to receipt. One could even imagine real-time interactions by way of the Bundle Protocol if the underlying network or communication resources were, in fact, capable of such communication. The logic behind the DTN Bundle Protocol was simply that real-time responses should not be *expected* although they could be accepted.

The Bundle Protocol provides a store-and-forward message service much like the Internet Protocol, except the Bundle Protocol allows for arbitrarily long storage at intermediate nodes. Because existing transport protocols would not accommodate such delays, the Bundle Protocol runs on top of transport protocols in networks that support them, such as the Internet. In other networks where it wouldn't make sense to use IP or IP routing, the Bundle Protocols can run directly over data link layers. Examples of networks that might not use IP include the Deep Space Network as well as low-power sensor networks or locally-connected but highly mobile networks.

Another important attribute of the Bundle Protocol is the concept of *custody transfer* used for reliable transmission. This notion was explored in an earlier design by the JPL and MITRE teams called the CCSDS File Delivery Protocol (CFDP). In both CFDP and the Bundle Protocol, the current custodian retains a copy of the data (file or bundle) until a subsequent relay confirms that it is taking custody. This allows the point of retransmission to move forward in the network with each new custodian unlike the TCP protocol that uses end-to-end retransmission to recover from data loss. Thus if custody is successfully transferred from Earth to a Mars orbiter and the file or bundle is lost in transit from the orbiter to the surface, it can be retransmitted from the orbiter, and not from Earth. In the Bundle Protocol, taking custody of a bundle is a local decision. While it is desirable for relays to take custody of bundles that request it, certain relays may forward bundles without taking custody of them, and some may even serve as forwarders only and be unable take custody of bundles. In these cases, the bundles continue on their path until some later relay can take custody of them.

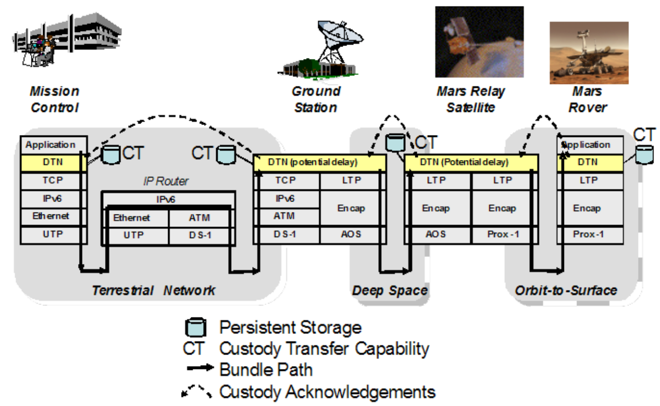

Figure 1: An end-to-end DTN transfer with custody.

Figure 1 shows how DTN might function in an end-to-end communications path from a mission control center to a rover on the surface of Mars. In the wired terrestrial Internet where connectivity is reliable and delays are low, DTN can run over TCP/IP, and may traverse a large piece of the Internet between DTN routers. At the ground station, DTN may choose to use a different mechanism such as the Licklider Transport Protocol (LTP [4-6]) over the Advanced Orbiting Systems (AOS) data link to communicate with a Mars relay satellite. In the last hop, the orbiter might use LTP over the Proximity-1 data link protocol to communicate with the rover. The switches from TCP/IP to LTP/AOS, to LTP/Prox-1 highlight the ability of DTN to choose transport layers appropriate to the local environment and to change mechanisms along the path. Custody transfers are illustrated via the storage devices marked ‘CT,’ with dashed lines indicating the acknowledgements that a downstream node has taken responsibility for ensuring data delivery.

A. Time Synchronization

One of the early realizations during the design of the Bundle Protocol was that the various nodes in an Interplanetary Internet would need at least loose time synchronization because celestial motions (orbits, planet and moon rotation) or other scheduling conventions would play a role in determining when nodes would be able to communicate. Nodes would need to have a common notion of time to support scheduled communications and correct antenna-pointing, for example. DTN leverages this implied knowledge to ensure that Bundles do not clog the system by marking each bundle with an explicit ‘time to live’ measured in seconds. Bundles in the network past their expiration times are removed from the network, and a notification may be sent to the source. One of the current areas of DTN research is investigating whether this mechanism alone is enough, or whether it should be augmented or replaced with a countdown timer. A countdown timer would allow bundles to be expired without requiring even loose clock synchronization, but would complicate some implementations that might want to consider bundles on removable memory devices as ‘in transit.’

B. Routing

This is a complex topic but in simple terms, the system needs to have some notion of network topology, probably computed at each node from local information about the node’s and its neighbors’ ephemerides. Conventional Internet routing protocols are not expected to work well under high delay and uncertain connectivity and it is expected that serious research will be required to develop alternatives. In the simplest scenarios, relatively fixed routing, together with suitable time synchronization and knowledge of communication opportunities may suffice for the interplanetary components of the network, while more conventional routing protocols and/or neighbor discovery protocols may suffice for low delay environments. Some experience has been gained with tests in tactical mobile scenarios that emulate some surface deployments one might imagine for local and possibly mobile sensor networks.

C. Security

From the beginning of the effort, security has been a high priority in the design of the Bundle Protocol. In addition to the common security goals of ensuring end-to-end data integrity and confidentiality, the DTN security protocols include a mechanism to protect the network itself from unwanted traffic. The rationale for this is that some resources in the Interplanetary Internet, especially bandwidth off-planet, are extremely precious. If someone were to inject bogus traffic, even if it were later rejected by the destination, the damage is done once that traffic uses one of the constrained resources such as the Deep Space Network. Contrast this with IP Security (IPSec), where false traffic injected into the network consumes network resources until it reaches the security destination. To combat denial-of-service attacks of this type, the DTN security protocols include a hop-by-hop security feature whereby a DTN router can drop incoming traffic that cannot be authenticated as coming from one of the router's trusted peers. This hop-by-hop security can use standard digital signature and certificate information commonly used in terrestrial networks.

D. Naming

As the design team considered the existing Domain Name System of the Internet[7] and its mechanisms, one aspect immediately stood out. In the DNS, communication is accomplished by translating a domain name such as www.jpl.nasa.gov into an Internet address such as 137.78.99.24 representing a 32 bit routable destination in the Internet. This translation requires access to a distributed network of domain name servers, and no communication is possible until the domain name is translated into an address. In the Interplanetary Internet, delays can be extreme and if one of the required domain name servers were off planet, it would drastically reduce communications efficiency by adding considerably to delay and also to potentially incorrect results. If the name/address combination is not fixed but varies as a result of destination mobility and changing connectivity to the Internet, the originating site could well receive out-of-date information as to the correct destination address and thus attempt to send a bundle to the wrong target. Consequently, the design team concluded that a form of delayed binding needed to be applied so that bundles could be routed towards the appropriate destination and only when the bundle arrives in the proper context, would the final destination address be determined and used for routing.

The naming conventions adopted for the Bundle Protocol are based on the Uniform Resource Identifier formats of the World Wide Web[8,9]. These conventions allow for a rich naming space, capable of accommodating complex mission scenarios in which many destinations are involved. While a rudimentary naming scheme has been in use for a number of years, research into exploiting the full power of URI-based naming and addressing is only now underway. The most basic scheme uses endpoint identifiers of the form dtn://XXX/YYY, where XXX identifies a particular bundle protocol relay and YYY identifies a service associated with that relay, much the way an IP address and port number identify an application in the Internet. An alternate approach under consideration is to encode both identity and location information separately in the XXX portion of the string and to allow the location information to be re-bound as the bundle traverses the network. This could support the notions of delayed binding and address aggregation, where relays on Mars could maintain a single routing entry for all bundles with 'Earth' as their destinations. The identity information could then be used *once the bundle reaches Earth* to re-bind the location to the IP address of the destination Bundle Protocol agent. Other researchers are interested in naming conventions useful in sensor networks, where entire database queries might be encoded in URI so that a bundle might be addressed to all sensor nodes with current temperature readings greater than 40 degrees centigrade.

E. Data Accountability

It may be important to account for the current positions of bundles within the network. For example, it might be important to know if a particular command had been radiated out of an antenna yet or not, or which node is currently storing or has custody of a particular bundle. The Bundle Protocol contains diagnostic capabilities to support all of these requirements. Data sources can request the various services and indicate the destination (not necessarily themselves) to which the diagnostic reports should be sent.

F. Applications

The primary objective of the Interplanetary Internet is to support manned and robotic exploration of the Solar System. Exchange of command and control information as well as transport of telemetry data will be the principal applications. In manned environments, one also expects the need to support streaming voice and video communication that would manifest itself as a sequence of bundles whose content can be interpreted as sound or moving imagery. It should be clear that in scenarios involving substantial distance, such as Earth to Mars, the notion of real-time interaction is simply untenable owing to the speed of light propagation delays. Even if continuous connectivity could be maintained, the information content would be delayed at least by 3 to 20 minutes depending on the positions of Earth and Mars in their respective orbits. There is probably some breakpoint in end-to-end latency and rate of disruptions where 'real-time' voice and video communication becomes unusable and people will prefer to substitute an asynchronous, email-like communication style.

While it is designed to *accommodate* long delays and intermittent connectivity, there is no reason the Bundle Protocol should *impose* those conditions if the network is connected and low-delay. Thus in the same way that UDP can support 'real-time' voice and video communications, the Bundle Protocol and its implementations should be able to support streaming applications when delays are low and connectivity is constant. In such networks, rapid bundle exchange might permit remote operation of mobile robots and management of manipulators via telepresence, for example. In terrestrial testing we have also found useful applications of the DTN protocols in sensor network design and in civilian and tactical mobile communication. Consequently, we anticipate that the Bundle Protocol may find utility in daily use on Earth as well as in space.

G. Scenarios

It is expected that early mission architectures will be relatively simple but could easily involve the need for store-and-forward Bundle Protocol operation. For example the Mars Phoenix lander is, as of this writing, relaying data it is collecting through a satellite that is programmed to receive, store, and forward data to the Deep Space Network, forming a very simple three point network (Earth, Mars and the relaying satellite). More complex missions involving multiple spacecraft requiring local communication, such as tandem formation flying for interferometric experiments, could readily require inter-spacecraft communication independent of command and control exchanges from Earth. The motivation of the design team is to craft a standard for space communication that is capable of accommodating highly delayed and disrupted environments while at the same time providing 'Internet-like' performance when the network is connected.

It is anticipated that on board spacecraft, conventional Internet TCP/IP protocols will suffice, as may the same protocols for local communications on the surfaces of other planets. However, between planets and in areas where disruption is common, the Bundle Protocol would be used to accommodate the delays and disruptions.

H. Testing

As this document is being written, plans have been underway to upload an implementation of the Bundle Protocol to the Deep Impact spacecraft now in orbit around the Sun, with testing to begin in October 2008. In 2009, it is planned to test the Bundle Protocol on board the International Space Station using equipment already on board and associated with the BIONET[10] project at the University of Colorado. Once these tests are completed, it is hoped that the Bundle Protocol and its implementations can be certified to be operating at NASA Technology Readiness Level 8 indicating readiness for these protocols to be used in live missions.

Tests of the Bundle Protocol have also been conducted with the support of the U.S. Defense Advanced Research Projects Agency in coordination with the U.S. Marine Corps and have produced results confirming the potential for the Bundle Protocol to deliver significantly more data than conventional TCP/IP in mobile tactical networks.

I. Standardization

Standardization of the Bundle Protocol by the international space community will enable interoperability between spacecraft from different agencies. Interoperability, coupled with cross-support agreements, should provide for increased data return, reduced mission risk, and reduced size, weight, and power requirements on spacecraft. Data return can be increased since there will be more communications opportunities; mission risk is reduced if there are multiple relay assets in the target area for a spacecraft and if the spacecraft can use those relay assets to get its data back to Earth. Finally, size, weight, and power can be reduced (for some spacecraft, such as landers) if they can transmit to relatively close orbiters instead of having to transmit all the way to Earth.

As more and more missions adopt the Bundle Protocol, each becomes a potential relay for future missions. The nature of cross-support agreements to take advantage of the interoperability provided by the protocol have yet to be worked out, but tests between NASA rovers and ESA orbiters using the Proximity-1 protocol have demonstrated that in-space cross-support is possible.

III. Conclusions

This brief essay has outlined the logic behind the development of a Delay and Disruption Tolerant Network design for an Interplanetary Internet. The design team hopes that successful testing of the Bundle Protocol will produce adoption as a standard and subsequent use by all space-faring nations in their future missions. Such an outcome would render all mission spacecraft technically interoperable, increasing the potential for re-purposing of existing mission assets to support subsequent missions. In effect, an Interplanetary backbone could be accreted over a period of years to decades, as a consequence of cumulative mission deployments.

Acknowledgments

The author acknowledges the long-standing and continuing contributions of colleagues at the Jet Propulsion Laboratory (Scott Burleigh, Leigh Torgerson, Wallace Tai, Jay Wyatt, Les Deutsch, Peter Shames, Loren Clare), Intel Berkeley Research Laboratory (Kevin Fall, Mike Demmer), The MITRE Corporation (Robert Durst, Keith Scott, Susan Symington), Sparta (Howard Weiss, Peter Lovell), NASA (Adrian Hooke, John Rush), Goddard Space Flight Center (David Israel, Jane Marquart), and many others too numerous to mention here. Special thanks to Keith Scott for considerable assistance with editing the description and illustration of the DTN architecture and Bundle Protocol design. This work was carried out under contract to the US Defense Advanced Research Projects Agency and the US National Air and Space Administration.

References

[1]Farrell, Stephen and Cahill, Vinny, *Delay- and Disruption-Tolerant Networking,* Artech House, Canton, MA, 2006 [ISBN-10 1596930632; ISBN-13 978-1596930636]

[2]Cerf, V. et al, "Delay-Tolerant Network Architecture," Request for Comment 4838, Internet Society, April 2007.

[3]Scott, K. and Burleigh, S, "Bundle Protocol Specification," Request for Comment 5050, Internet Society, November 2007.

[4]Burleigh, S; Ramadas, M; Farrell, S; "Licklider Transmission Protocol – Motivation", Request for Comment 5325, September 2008.

[5]Ramadas, M; Burleigh, S; Farrell, S; "Licklider Transmission Protocol – Specification," Request for Comment 5326, Internet Society, September 2008

[6]Farrell, S; Ramadas, M; Burleigh, S; Licklider Transmission Protocol – Security Extensions," Request for Comment 5327, Internet Society, September 2008

[7]See http://www.dns.net/dnsrd/rfc/ for a substantial bibliography of relevant documents on the DNS system.

[8]See http://en.wikipedia.org/wiki/Uniform_Resource_Identifier

[9]Berners-Lee, T., Fielding, R. and Masinter, L., "Uniform Resource Identifier (URI): Generic Syntax, RFC 3986, Internet Society, January 2005.

[10]http://www.colorado.edu/engineering/BioServe/portfolio.html